Substantive Bias in Phonotactic Learning:
Positional Extension of an Obstruent Voicing Contrast

Eleanor Glewwe
University of California, Los Angeles

A number of studies have used artificial grammar learning (AGL) to investigate the extent to which the phonological typology is shaped by synchronic learning biases (Wilson 2006, Moreton 2008, Finley 2012, White 2013, a.o.). Moreton & Pater's (2012) review concluded that there is robust evidence for complexity bias (bias toward featurally simpler patterns) but scant evidence for substantive bias (bias toward phonetically natural patterns). I present an experiment that examined substantive bias in phonotactic learning by testing whether subjects reproduced an attested and phonetically motivated phonotactic implicational. The results do not support a substantive bias account but suggest a possible effect of complexity bias.

The implicational I tested is this: if a language contrasts voicing in obstruents word-finally, it contrasts voicing in obstruents word-initially, but not necessarily vice versa (Steriade 1997). Using an AGL paradigm, I exposed subjects to an obstruent voicing contrast in either word-initial or word-final position and tested whether they extended the contrast to the other position. There were four training languages defined on two dimensions: Trained Contrast Position (whether the language exhibited an obstruent voicing contrast word-initially or word-finally) and Trained Neutralization Value (whether obstruents in the other position "neutralized" to voiced (D) or voiceless (T)). The DFinalContrast and TFinalContrast conditions featured an obstruent voicing contrast word-finally. In the DFinalContrast condition, subjects were exposed to items ending with both voiceless stops (T#) and voiced stops (D#) and items beginning with voiced stops (#D), but not items beginning with voiceless stops (#T). Subjects in the TFinalContrast condition heard T#, D#, and #T items, but not #D items. The InitialContrastD and InitialContrastT conditions featured an obstruent voicing contrast word-initially. Subjects in the InitialContrastD condition heard #T, #D, and D# items while subjects in the InitialContrastT condition heard #T, #D, and T# items. Items were of the form CVCVC, and the other two Cs in each item were sonorants, so subjects in all four conditions heard word-initial and word-final sonorants.

In the training phase, subjects were told they would be learning words of a new language. Each word was paired with an image. In the test phase, subjects heard additional words, without images, and had to say whether the word could be a word of the language they had just been listening to or not. The test items were the same in all four conditions and consisted of #T, #D, T#, and D# items. Test items fell into three categories: familiar conforming (items heard in training), novel conforming (new items consistent with the pattern heard in training), and novel nonconforming (new items of the type not heard in training; e.g. D# in the InitialContrastT condition).

Subjects' acceptance rates of novel nonconforming items, relative to novel conforming items, indicate whether they have extended the obstruent voicing contrast to a new position in a given condition. If learners are biased toward phonetically natural phonotactic systems, subjects exposed to the obstruent voicing contrast word-initially should less readily accept novel nonconforming items than subjects exposed to the obstruent voicing contrast word-finally. In terms of the effect of Trained Neutralization Value, I predicted that if there was any effect, subjects would extend from voiced to voiceless obstruents more than from voiceless to voiced obstruents, since voiced obstruents are more marked. If learners are not biased toward

phonetically natural phonotactic systems, subjects should accept novel nonconforming items at similar rates across all four conditions, since all four training patterns are of equal formal complexity and should therefore be equally learnable.

Figure 1 shows the mean acceptance rate across subjects of novel conforming and novel nonconforming items in each of the four conditions. Acceptance rates of novel conforming items are above chance in all four conditions and do not differ significantly between conditions. I fit a mixed-effects logistic regression to the novel nonconforming items with response (accept or reject) as the dependent variable and Condition as the fixed effect and then carried out post-hoc pairwise comparisons between conditions. Subjects did not demonstrate greater extension of the voicing contrast from word-final to word-initial position than from word-initial to word-final position. Although the difference in acceptance rate of novel nonconforming items between the DFinalContrast condition and the InitialContrastD condition was in the direction that would support a positional substantive bias, it was not significant. The difference in acceptance rate between TFinalContrast and InitialContrastT was in the *opposite* direction and was also not significant. The acceptance rate of novel nonconforming items in InitialContrastD was significantly lower than the acceptance rate in InitialContrastT. This runs counter to the predicted effect of Trained Neutralization Value, so subjects did not extend more from voiced to voiceless obstruents than from voiceless to voiced obstruents. The results therefore do not yield evidence for any substantive biases.

Fig. 1: Acceptance Rates of Novel Items by Condition

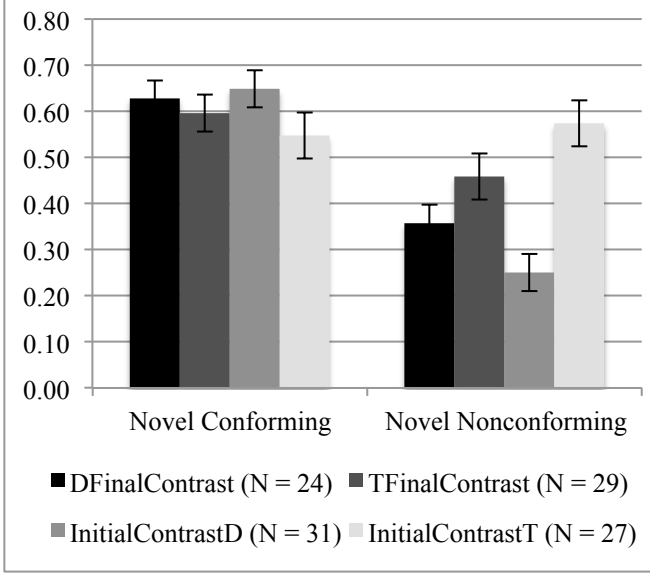

Why were InitialContrastT subjects more willing to accept nonconforming D# items than InitialContrastD subjects were to accept nonconforming T# items? A possible explanation for the difference between these conditions lies with the presence of word-initial and word-final sonorants in the stimuli. InitialContrastD subjects could have posited a constraint *[-voice]# to exclude the type of item they were not exposed to (T#) while InitialContrastT subjects would have had to posit the more complex constraint *[-son, +voice]# to exclude D# items. The simpler constraint may have been easier to learn, resulting in InitialContrastD subjects' superior accuracy at rejecting nonconforming items relative to InitialContrastT subjects. Thus the results suggest a possible role of complexity bias in phonotactic learning. The same difference in constraint complexity applies to the DFinalContrast and TFinalContrast conditions, and indeed the difference in acceptance rate of novel nonconforming items between these two conditions was in the direction that would support an effect of complexity bias. However, the difference in acceptance rate was not significant.

This study adds to the small body of work explicitly testing substantive bias in phonotactic learning. It further calls into question the influence of substantive bias on phonological learning while leaving room for potential effects of complexity bias.